

996 4419 DE



PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION  
International Bureau


## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : <b>H04L 12/44</b>		A1	(11) International Publication Number: <b>WO 97/29573</b>
			(43) International Publication Date: 14 August 1997 (14.08.97)
(21) International Application Number: PCT/US96/14197		(81) Designated States: AL, AM, AT, AT (Utility model), AU, AZ, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), TJ, TM, TR, TT, UA, UG, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).	
(22) International Filing Date: 4 September 1996 (04.09.96)			
(30) Priority Data: 08/600,414      9 February 1996 (09.02.96)      US			
(71) Applicant: LEVEL ONE COMMUNICATIONS, INC. [US/US]; 9750 Goethe Road, Sacramento, CA 95827 (US).			
(72) Inventors: FEUERSTRAETER, Mark, T.; 3441 Data Drive #537, Rancho Cordova, CA 95670 (US). KUBES, Steven; 1210 Shetland Way, El Dorado Hills, CA 95672 (US).			
(74) Agent: GATES, George, H.; Merchant, Gould, Smith, Edell, Welter & Schmidt, Suite 400, 11150 Santa Monica Boulevard, Los Angeles, CA 90025 (US).		Published With international search report.	

(54) Title: AUTOMATIC SPEED SWITCHING REPEATER

## (57) Abstract

A network repeater (300) having automatic speed switching capability. The repeater includes a first repeater logic means (306) for connecting devices operating at the first rate to a first backplane (310), a second repeater logic means (308) for connecting devices operating at the second rate to a second backplane (312), and a port switching fabric (304), coupled to the first and second repeater logic means, for determining the transmission rate of a signal at a port (302) and routing the signal to one of the repeater logic according to the transmission rate determination. The repeater includes a serial controller for accessing internal management counters, LEDs for providing status information about the repeater, and a media access controller for controlling functions including the transmission, loopback and reception of Ethernet frames.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

## AUTOMATIC SPEED SWITCHING REPEATER

### BACKGROUND OF THE INVENTION

#### 1. Field of the Invention.

This invention relates in general to a network device, and in particular, to a network repeater having automatic speed switching capability.

#### 5 2. Description of Related Art.

Recent advancements in the art of data communications have provided great strides in resource sharing amongst computer systems through the use of networks which offer reliable high-speed data channels. Networks allow versatility by defining a common standard for communication so that information independent of vendor equipment may be exchanged across user applications. As the popularity of networks increase so does the demand for performance. More sophisticated protocols are being established to meet this demand and are utilizing existing twisted pair wires in office buildings so that virtually all computer literate users have access to resources with minimal expense.

15 As will be appreciated by those skilled in the art, communication networks and their operations can be described according to the Open Systems Interconnection (OSI) model which includes seven layers including an application, presentation, session, transport, network, link, and physical layer. The OSI model was developed by the International Organization for Standardization (ISO) and is described in "The Basics Book of OSI and Network Management" by Motorola Codex from Addison-Wesley Publishing Company, Inc., 1993 (First Printing September 1992).

Each layer of the OSI model performs a specific data communications task, a service to and for the layer that precedes it (e.g., the network layer provides a service for the transport layer). The process can be likened to placing a letter in a series of envelopes before it is sent through the postal system. Each succeeding envelope adds another layer of processing or overhead information necessary to process the transaction. Together, all the envelopes help make sure the letter gets to the right address and that the message received is identical to

the message sent. Once the entire package is received at its destination, the envelopes are opened one by one until the letter itself emerges exactly as written.

In a data communication transaction, however, each end user is unaware of the envelopes, which perform their functions transparently. For example, an automatic bank teller transaction can be tracked through the multilayer OSI system. One multiple layer system (Open System A) provides an application layer that is an interface to a person attempting a transaction, while the other multiple layer system (Open System B) provides an application layer that interfaces with applications software in a bank's host computer. The corresponding layers in Open Systems A and B are called peer layers and communicate through peer protocols. These peer protocols provide communication support for a users application, performing transaction related tasks such as debiting an account, dispensing currency, or crediting an account.

Actual data flow between the two open systems (Open System A and Open System B), however, is from top to bottom in one open system (Open System A, the source), across the communications line, and then from bottom to top in the other open system (Open System B, the destination). Each time that user application data passes downward from one layer to the next layer in the same system more processing information is added. When that information is removed and processed by the peer layer in the other system, it causes various tasks (error correction, flow control, etc.) to be performed.

The ISO has specifically defined all seven layers, which are summarized below in the order in which the data actually flows as they leave the source:

Layer 7, the application layer, provides for a user application (such as getting money from an automatic bank teller machine) to interface with the OSI application layer. That OSI application layer has a corresponding peer layer in the other open system, the bank's host computer.

Layer 6, the presentation layer, makes sure the user information (a request for \$50 in cash to be debited from your checking account) is in a format (i.e., syntax or sequence of ones and zeros) the destination open system can understand.

Layer 5, the session layer, provides synchronization control of data between the open systems (i.e., makes sure the bit configurations that pass through layer 5 at the source are the same as those that pass through layer 5 at the destination).

- 5        Layer 4, the transport layer, ensures that an end-to-end connection has been established between the two open systems and is often reliable (i.e., layer 4 at the destination confirms the request for a connection, so to speak, that it has received from layer 4 at the source).

- 10       Layer 3, the network layer, provides routing and relaying of data through the network (among other things, at layer 3 on the outbound side an address gets slapped on the envelope which is then read by layer 3 at the destination).

Layer 2, the data link layer, includes flow control of data as messages pass down through this layer in one open system and up through the peer layer in the other open system.

- 15       Layer 1, the physical interface layer, includes the ways in which data communications equipment is connected mechanically and electrically, and the means by which the data moves across those physical connections from layer 1 at the source to layer 1 at the destination.

- 20       The primary standard for Local and Metropolitan Area Network technologies is governed by IEEE Std. 802. IEEE Std. 802 describes the relationship among the family of 802 standards and their relationship to the ISO OSI Basic Reference Model. Generally, IEEE Std. 802 prescribes the functional, electrical and mechanical protocols, and the physical and data link layers for Local and Metropolitan Area Networks (LAN/MAN). The specification
- 25       augments network principles, conforming to the ISO seven-layer model for OSI, commonly referred to as "Ethernet". In the hierarchy of the seven-layer model, the lowest layers, the so-called physical and data link layers, comprise functional modules that specify the physical transmission media and the way network nodes interface to it, the mechanics of transmitting information over the media in an
- 30       error-free manner, and the format the information must take in order to be transmitted.

While there are several LAN technologies in use today, Ethernet is by far the most popular. The definitions of an Ethernet Repeater functions are contained in the IEEE 802.3 specification, which is herein incorporated by reference. This standard defines attributes which can be used by a management  
5 function within an IEEE 802.3 Ethernet Repeater to monitor network behavior.

The vast majority of computer vendors today equip their products with 10 Mbps Ethernet attachments, making it possible to link all manner of computers with an Ethernet LAN. Nevertheless, the need for faster data transmission has led to the development of Fast Ethernet standards which carry Ethernet frames at  
10 100 Mbps. When the IEEE standardization committee began work on a faster Ethernet system, two approaches were presented. One approach was to speed up the original Ethernet system to 100-Mbps, keeping the original CSMA/CD medium access control mechanism. This approach is called 100BASE-T Fast Ethernet.

15 Another approach presented to the committee was to create an entirely new medium access control mechanism, one based on hubs that controlled access to the medium using a "demand priority" mechanism. This new access control system transports standard Ethernet frames, but it does it with a new medium access control mechanism. This system was further extended to allow it to  
20 transport token ring frames as well. As a result, this approach is now called 100VG-AnyLAN.

The IEEE decided to create standards for both approaches. The 100BASE-T Fast Ethernet standard described here is part of the original 802.3 standard, which was incorporated by reference earlier. The 100VG-AnyLAN  
25 system is standardized under a new number: IEEE 802.12, herein incorporated by reference.

The Fast Ethernet standards both include mechanisms for Auto-Negotiation of the media speed. As the 100 Mbps standard becomes more widely adopted, computers are being equipped with an Ethernet interface that  
30 operates at both 10 Mbps and 100 Mbps. The Auto-Negotiation function is an optional part of the Ethernet standard that makes it possible for devices to exchange information about their

abilities over a link segment. This, in turn, allows the devices to perform automatic configuration to achieve the best possible mode of operation over a link. At a minimum, Auto-Negotiation can provide automatic speed matching for multi-speed devices at each end of a link. Multi-speed Ethernet interfaces  
5 can then take advantage of the highest speed offered by a multi-speed hub port.

The Auto-Negotiation protocol includes automatic sensing for other capabilities as well. For example, a switched hub that is capable of supporting full duplex operation on some or all of its ports can advertise that fact with the Auto-Negotiation protocol. Interfaces connected to the switch hub port that also  
10 support full duplex operation can then configure themselves to use the full duplex mode in interaction with the hub.

Integrated multi-port repeaters exist which provide separately either 10 Mb/s or 100 Mb/s functionality. However, the incorporation of 10 Mb/s or 100 Mb/s repeater functionality within a single repeater is beyond the scope of IEEE  
15 802.3 standard.

Fig. 9 illustrates the traditional solution 900 of providing 100 Mbps functionality in a network system along with 10 Mbps functionality. A 10 Mbps repeater 910 is connected to a 100 Mbps repeater 912 through a bridge 914. Ten Mbps devices 916 are connected to the 10 Mbps repeater 910. One hundred  
20 Mbps devices 918 are connected to the 100 Mbps repeater 914. Upgrading a device to Fast Ethernet requires disconnecting the device 916 from the 10 Mbps repeater 910 and re-routing cable from the device 916 to the 100 Mbps repeater 914. However, this procedure is time consuming and confusing. Furthermore, to provide the dual speed functionality, both types of repeaters and a bridge must  
25 be purchased. Thus, the traditional solution is too expensive.

If a 10 Mbps segment is replaced with a 100 Mbps segment, fast link pulses (FLPs) will be received by the repeater hub, and the Auto-Negotiation protocol will result in the hub port operating at 100-Mbps as long as all interfaces connected to the repeater hub can operate at 100-Mbps. The change  
30 from 10-Mbps to 100-Mbps will occur with no manual intervention.

Auto-Negotiation ensures that all devices attached to the hub are operating at the highest common denominator. Since a repeater hub is used to create a

shared signal channel for all devices attached to the repeater ports, that shared signal channel must operate no faster than the slowest device attached to it.

If an repeater hub has one of its ports attached to a device that only supports 10 Mbps transmissions while the rest of the ports are attached to 100 Mbps devices, the hub will negotiate a speed of 10-Mbps for all ports, since that is the highest common denominator for all repeated ports. When every device attached to the repeater hub is capable of operating at 100-Mbps, then the hub will negotiate 100-Mbps for all ports.

If there is no common technology detected at either end of the link, then the Auto-Negotiation protocol will not make a connection, and the port will be left in the off condition.

It can therefore be seen that there is a need for a repeater which provides both 10 Mb/s or 100 Mb/s functionality.

It can also be seen that there is a need for a repeater which automatically switches to either a 10 Mb/s backplane/segment or a 100 Mb/s backplane/segment.

### SUMMARY OF THE INVENTION

To overcome the limitations in the prior art described above, and to overcome other limitations that will become apparent upon reading and understanding the present specification, the present invention discloses a network repeater having automatic speed switching capability.

The present invention solves the above-described problems by providing a first repeater logic means for connecting devices operating at the first rate to a first backplane, a second repeater logic means for connecting devices operating at a second rate to a second backplane, and a port switching fabric, coupled to the first and second repeater logic means, for determining the transmission rate of a signal at a port and routing the signal to one of the repeater logic according to the transmission rate determination.

One aspect of the present invention is that the repeater includes an integrated serial controller for accessing internal management counters, control and status information.



Another aspect of the present invention is that the repeater includes LEDs for providing status information about the repeater.

Yet another aspect of the present invention is that the repeater includes a media access controller for controlling functions including the transmission,  
5 loopback and reception of Ethernet frames.

Another aspect of the present invention is that the repeater includes a direct memory access controller for servicing the media access controller and controlling memory.

Still another aspect of the present invention is that a direct memory access  
10 controller transfers data from the receive FIFO to memory during a reception, and from memory to a transmit FIFO during transmission.

Another aspect of the present invention is that a dual buffer ring structure manages the transmit and receive traffic.

Another aspect of the present invention is that a link list structure could  
15 be used to chain frames between repeaters of different transmission rates on different backplanes.

Another aspect of the present invention is that a buffer area is disposed between the two rings in memory to prevent packet fragmentation so that an arriving packet that does not fit within the space allotted between the buffer start  
20 address and the ring stop boundary is written past the ring stop boundary until the data buffer is complete and the start of the next frame will be wrapped to the ring's start address.

Still another aspect of the present invention is that the repeater includes bridging means having a receiving and forwarding state, the receiving state for  
25 receiving signals from remote media access control managers and the forwarding state for responding to remote media access control managers.

These and various other advantages and features of novelty which characterize the invention are pointed out with particularity in the claims annexed hereto and form a part hereof. However, for a better understanding of the  
30 invention, its advantages, and the objects obtained by its use, reference should be made to the drawings which form a further part hereof, and to accompanying

descriptive matter, in which there is illustrated and described specific examples of an apparatus in accordance with the invention.

### BRIEF DESCRIPTION OF THE DRAWINGS

5 Referring now to the drawings in which like reference numbers represent corresponding parts throughout:

Fig. 1 illustrates an automatic speed switching repeater according to the present invention;

Fig. 2 illustrates an exemplary system architecture using an automatic speed switching repeater according to the present invention;

Fig. 3 illustrates the chip architecture of an automatic speed switching repeater according to the present invention;

Fig. 4 illustrates a second embodiment for a system architecture using an automatic speed switching repeater according to the present invention wherein an external bridge/manage card is provided;

Fig. 5 illustrates an alternative chip architecture for an automatic speed switching repeater according to the present invention;

Fig. 6 illustrates automatic bridging between a 10 Mbps media access controller and a 100 Mbps media access controller according to the present invention;

Fig. 7 illustrates a logical DMA channel according to the present invention;

Fig. 8 illustrates a ring managed, two-channel DMA bridge according to the present invention;

Fig. 9 illustrates a prior art solution for providing 100 Mbps functionality in a network system along with 10 Mbps functionality.

### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

In the following description of the exemplary embodiment, reference is made to the accompanying drawings which form a part hereof, and in which is shown by way of illustration a specific embodiment in which the invention may be practiced. It is to be understood that other embodiments may be utilized and

structural changes may be made without departing from the scope of the present invention.

The present invention provides a repeater which automatically switches to either a 10 Mbps backplane/segment or a 100 Mbps backplane/segment.

5 Fig. 1 illustrates a fully integrated IEEE 802.3 repeater incorporating both 10 and 100 Mbit/sec functionality. The repeater has several ports 12 which are connected to physical media segments 14 that make up the Ethernet system 16. Some of the media segments 18 provide twisted pair paths for transmission rates of 10 Mbps. Other media segments 20 provide a path to 100 Mbps devices.

10 When a device 22 is changed from 10 Mbps to 100 Mbps no wiring or altering of connections is necessary. The repeater 10 automatically provides functionality to the new 100 Mbps device 22.

Fig. 2 illustrates an exemplary system architecture 200 according to the present invention. In Figure 2, three repeaters 202, 204, 206, according to the invention, are connected to a 10 Mbps backplane 208 and a 100 Mbps backplane 210. Each repeater is connected to distinct media segments 212.

The repeater includes on-board, fully 802.3 compliant media access controller (MAC) entities 214, 222. These units perform the media access control functions for the transmission, loopback and reception of Ethernet frames. A MAC utilizes both a transmit and receive FIFO to provide latency for a DMA controller which services the MAC. During transmission, the MAC may take standard LLC frames (provided by the DMA) and serialize the data. The hardware may then append a preamble followed by a Start of Frame Delimiter, as a header to the outgoing data. In addition, the hardware can optionally  
20 append a computed Frame Check Sequence to the end of the outgoing frame for error detection purposes.

When a device operating at 10 Mbps is connected to a repeater 206, and that device's signal is to be repeated and routed to a 100 Mbps device connected to a port on repeater 204, the media access controller 214 in 10 Mbps repeater  
30 206 routes a signal to a 10/100 bridge memory 220 which in turn bridges the signal to the media access controller 222 in 100 Mbps repeater 204.

Each of the media access controllers 214, 222 are connected to content addressable memory (CAM) 226. The content addressable memory 226 is a storage device which includes comparison logic with each bit of storage. A data value is broadcast to all words of storage and compared with the values there.

- 5 Words which match are flagged in some manner, and subsequent operations can then work on flagged words. Furthermore, the repeater may act as a media access controller management device 228. Thus, management commands may be routed from the repeater 202 to CPU 230.

Fig. 3 illustrates the chip architecture 300 of a repeater according to the present invention. The repeater contains ports 302 for connection to 10 Mbps  
10 twisted pair, 100 Mbps twisted pair or fiber optics (coax can be used at 10 Mbps connections). The ports 302 route the signals to port switching fabric 304 which provides the auto negotiation functionality.

The auto-negotiation functionality, as defined in IEEE Std 802.3u/D4-  
15 1995 (Draft supplement to ISO/IEC 8802-3:1993 ANSI/IEEE Std 802.3-1993 ed.), herein incorporated by reference, negotiates the transmission rate at each port and routes the signal to the appropriate repeater.

Auto-Negotiation is defined in Clause 28 of the D4 draft of the ANSI/IEEE Std 802.3 MAC Parameters, Physical Layer, Medium Attachment  
20 Units and Repeater for 100 Mbps Operation. The Auto-Negotiation program takes control of the cable when a connection is established to a network device and detects the various modes that exist in the device on the other end of the wire.

Auto-Negotiation acts like a rotary switch that automatically switches to  
25 the correct technology, such as 10BASE-T, 100BASE-TX, 100BASE-T4, or a corresponding Full Duplex mode. Once the performance mode for the device is determined, Auto-Negotiation passes control of the cable to the appropriate technology and becomes transparent until the connection is broken.

The primary benefit of Auto-Negotiation is the automatic connection of  
30 the appropriate performance technology without any intervention from a user, manager, or management software. As new nodes are connected which have 100 Mbps functionality, the port switching fabric of the present invention

determines whether the port is transmitting at 10 Mbps or 100 Mbps and routes the signal to the proper repeater logic 306, 308.

The basic mechanism that Auto-Negotiation uses to advertise a device's abilities is a series of link pulses which encode a 16 bit word, known as a Fast Link Pulse (FLP) Burst. An FLP Burst is composed of 17 to 33 link pulses which are identical to the link pulses used in 10BASE-T to determine whether a link has a valid connection (sometimes referred to as Normal Link Pulses or NLPs.) FLP Bursts occur at the same interval as NLPs, 16.8ms. An FLP Burst may have nominal duration of 2 ms.

10 An FLP Burst interleaves clock pulses with data pulses to encode a 16 bit word. The absence of a pulse within a time window following a clock pulse encodes a logic zero and a pulse within the time window following a clock pulse encodes a logic one.

The key to Auto-Negotiation's flexibility and expendability is the encoding of the 16 bit word. The 16 bit word is referred to as the Link Code Word (LCW). Auto-Negotiation must ensure that the node receives the Link Code Word correctly and that the node's Link Code Word is received correctly in order to make a connection decision. Auto-Negotiation uses the Arbitration function to accomplish this.

20 Once the port switching fabric 308 has ascertained the transmission rate it routes the Ethernet frame to the appropriate repeater logic 306, 308. 10 Mbps repeater logic 306 and 100 Mbps repeater logic 308 are coupled to the port switching fabric 304. The 10 Mbps repeater logic 306 routes a signal onto the 10 Mbps backplane 310. The 100 Mbps repeater logic 308 routes a signal to the 100 Mbps backplane 312. The repeater 300 also includes CPU/serial/management functionality 314 along with FIFO buffers 316. A 16 or 32 bit bus master direct memory access 318 is coupled to memory 320 which may be external to the repeater chip itself.

30 Since the repeater uses the CSMA/CD access method, collisions are a fact of life. In the event of a transmit collision, the repeater will wait a random amount of time, as determined by the Truncated Binary Exponential Backoff Algorithm, before retransmitting the frame.

The repeater 300 takes the serial bit stream coming from the port switching fabric 304, strips off the preamble/SFD header, and deserializes the data into word quantities. This data passes into the receive FIFO 316 where it awaits service from its supporting DMA channel 318. At the end of reception, 5 the repeater checks for any errors that may have occurred (such as CRC and length errors) and may provide the user with the ability to filter the errored packet.

As mentioned above, the repeater contains an on-board DMA macro 318 which transfers data from a receive FIFO 316 to memory 320 during a reception, 10 and from memory 320 to a transmit FIFO 322 during transmission. All repeater bus master activity takes place in block sizes which can be programmed by the user.

The repeater uses a buffer ring structure to manage both the transmit and receive traffic. Two static registers and two dynamic registers govern the 15 smooth operation of a ring. The amount of address space dedicated to a particular ring structure may be defined by the Ring Start (RSTRT) and Ring Stop (RSTP) registers. These pointers are setup by the software during initialization. The other two registers define a Current Pointer (CURR) and a Buffer Start Pointer (BSTRT). During reception, the CURR indicates where the 20 receive DMA will write its next word transfer. The BSTRT contains the address pointing to the beginning of the current frame being received. This allows the DMA to quickly recover its pointers should the MAC reject the incoming data. During transmission, the CURR indicates where the transmit DMA is currently reading from (or writing to) in memory. Similarly, the BSTRT is a pointer to 25 the starting address of the current transmit buffer. This provides an easy mechanism for retransmitting frames that experience collisions.

In one embodiment of the present invention two restraints are associated with management of system memory. A first constraint involves the programming of certain registers within the repeater. First, ring boundaries are 30 set preferably on 2048 byte pages. Second, data buffers start preferably on 8 byte boundaries. Thus, the repeater's addressable memory is broken up into

sixteen 64 kbyte blocks. The user may select which block will be accessed by programming the page registers.

The second constraint involves the lack of packet fragmentation within a ring buffer. Space should be left between any two rings in memory. The amount of space should be large enough to encompass the interface's MTU size plus descriptor (2 kbytes) plus an additional 2kbytes. This is needed for the cases when an arriving packet will not fit within the space allotted, between the buffer start address and the ring stop boundary. Rather than have the hardware fragment the packet by wrapping the overflow data back around to the beginning of the ring, the packet will instead be written past the ring stop boundary until the data buffer is complete. The start of the next frame will be wrapped to the ring's start address. This eases the processing burden of software since it is much easier to deal with the non-fragmented blocks.

Fig. 4 illustrates an alternative systems architecture embodiment wherein a bridge/manage card 402 is provided external to the hub 404. Each repeater 406, 408, 410 within the hub is programmed to provide functionality for the two backplanes 412, 414. The repeater routes the signal to an appropriate media access controller 416, 418 in the bridge/manage card. The 10 MPBS media access controller 416 communicates with the 100 media access controller 418 using the bridging technique discussed earlier. The CPU 420, content addressable memory 422, the serial controller 424, and an additional memory 426 for storing packets are also provided in the bridge manage card 402.

Fig. 5 illustrates an alternative chip architecture 500 according to the present invention. The repeater contains ports 502 for connection to 10 Mbps twisted pair, 100 Mbps twisted pair or fiber optics. The ports 502 route the signals to port switching fabric 504 which provides the auto negotiation functionality.

Once the port switching fabric 308 has ascertained the transmission rate it routes the Ethernet frame to the appropriate repeater logic 306, 308. 10 Mbps repeater logic 306 and 100 Mbps repeater logic 308 are coupled to the port switching fabric 304. The 10 Mbps repeater logic 306 routes a signal onto the 10 Mbps backplane 310. The 100 Mbps repeater logic 308 routes a signal to the

100 Mbps backplane 312. The repeater 300 also includes CPU/serial/management functionality 314 along with FIFO buffers 316. A 16 or 32 bit bus master direct memory access 318 is coupled to memory 320 which may be external to the repeater chip itself.

5           The repeater may also include a media independent interface 550. The medium independent interface 550 is a set of electronics that provides a way to link the Ethernet medium access control functions in the network device 500 with the physical layer device that sends signals onto the network medium. A medium independent interface 550 may optionally support both 10 Mbps and 100  
10   Mbps operation, allowing suitably equipped network devices to connect to both 10 base-T and 100 base-T twisted paramedia segments.

          The medium independent interface 550 is designed to make the signaling differences among the various media segments transparent to the Ethernet chips in the network device. The medium independent interface 550 converts the line  
15   signals received from the various media segments by the transceiver into digital format signals that are then provided to the Ethernet chips in the device. The repeater may also include LED drivers 560 for providing status communication to the user.

          Fig. 6 illustrates the automatic bridging 600 between a 10 Mbps media  
20   access controller 602 and a 100 Mbps media access controller 604. Signals are routed between the 10 Mbps media access controller 602 and the 100 Mbps media access controller 604 using memory rings 606, 608.

          As part of the Simple Network Management Protocol (SNMP), a supporting management device must be able to both receive and respond to a  
25   remote manager's request for statistical information. Accordingly, an on-board MAC is needed to process these management frames. To complete the repeater Hub Management system, a CPU running SNMP software must be provided. The CPU interprets the SNMP management frames and provides all requested information. This is accomplished by accessing the repeater statistical data base  
30   and passing all pertinent data into an SNMP frame, and then sending the packet back to the requesting manager. The DMA macro frees up more bus bandwidth, as well as simplifies any driver software.



While the QFR bridge is in its forwarding state, there are two possibilities modes of operation: (1) operation with CPU intervention; and (2) operation without CPU intervention.

5     **1.     Bridging using the CPU**

The MAC in the receiving port must lock on to the Start of Frame Delimiter of an 802.3 packet. Next, the nature of the current frame must be determined via its destination address. This information is given to the receive DMA channel so it can set up the correct pointers in memory.

10       After the number of bytes received has exceeded a programmable threshold level, a receive bus master will request service from a DMA. After the DMA buffers both data and descriptor information to memory, the bus master may interrupt the host to indicate that an Ethernet frame is ready to be forwarded. Upon receiving an interrupt, the CPU pushes its flags and instruction  
15     pointer onto its stack, access its interrupt table, and begin executing the correct Interrupt Service Routine (ISR). During this particular service routine, the CPU, acquires the bus and performs a read cycle to the receive DMA in order to determine where in memory the frame can be found. Once this buffer address has been determined, the CPU can begin to transfer data to a MAC-only slave  
20     device found at the bridge's other port(s). In order to complete the block data transfer, two options exist: (1) Use a DMA channel; or (2) Use CPU memory cycles.

**a.     Using a DMA Channel**

25       If an embedded processor (such as a 186) is used as the resident CPU, at least one internal DMA channel is available to the programmer. After the CPU gets an indication that a frame is ready to be forwarded, a DMA channel may be setup to take care of the block transfer.

30       **b.     Using CPU Memory Cycles**

The alternative to using a DMA macro is for the CPU to perform memory to register transfers, followed by a register to IO instruction. This

requires that the CPU first bring the data bound for the network into an internal register, then transfer that data to a MAC-only slave device.

## 2. Bridging with No CPU Intervention

5 Strictly speaking, the CPU will more than likely have some involvement in the forwarding process, albeit a very minimal one. The alternative to using a CPU in the forwarding process, is to instead use a device which has an on-board DMA macro that can "speak" the same protocol as the repeater's DMA. In this situation, the bus bandwidth utilization becomes much more efficient. A bus  
10 master's DMA is able to perform bus cycles at a rate that is only limited by its target device's response time, in this case, the SRAM's cycle time. Accordingly, the latter of the two approaches is preferable.

### a. QFR Bridging Using a Companion Device

15 The memory structure used by the repeater consists of a logical ring plus pointers. The primary responsibility of maintaining these ring's falls on the writing device, or the Receive Bus Master. The receive bus master will keep pointers which define the rings upper and lower bounds, the start of the current frame's buffer, the location of where the next frame will be stored, and the  
20 actual location of the current write pointer.

Within the bridge, the receiving bus master of one port will communicate with a transmitting bus master of another port. This communication between the ports may be considered a single logical DMA channel. A transmitting device on the other port will maintain a similar set of registers as that found in a receive  
25 bus master.

Fig. 7 shows a graphical representation 700 of a logical DMA channel 708 from a receive device in port A 702 to a transmit device in port B 704.

A bridge that is 802.1d compliant must be able to forward traffic in both directions, that is, from port A 702 to port B 704 and from port B 704 to port A  
30 702. The second data path can be supported through a second logical DMA channel. Fig. 8 shows a ring managed, two-channel DMA bridge 800.

In order to support a second DMA channel, another logical ring 810 must be created in memory. For the example shown in Fig. 8, this new ring structure 810 would be managed by the writing bus master, or port B's 804 receive device.

- 5           The main responsibility of an 802.1d compliant bridge is to forward traffic of unknown destination and to reject all others. However, since a MAC layer bridge contains the MAC entities needed to gain access to the particular working group supported (i.e. Ethernet, Token Ring, Token Bus...), the bridge then also has the ability to be uniquely addressed through the MAC entity
- 10           contained within the bridge. This implies that there are some frames which should not be forwarded to another LAN segment, yet neither should they be rejected. Such frames can be regarded as Host traffic to the bridge. That is, they were addressed to the bridge, and no one else.

- In the case of the repeater, the best example of Host traffic is that of
- 15           SNMP in-band management frames. The term in-band means that the remote manager uses the media on which it is gathering statistics to transfer the SNMP frames. When an SNMP packet arrives, it is not forwarded. Instead, it is routed to a region of memory that has been reserved for Host traffic. Once the DMA activity has stopped, the CPU is given a maskable interrupt indicating that
- 20           Host traffic has just arrived. As part of the descriptor field, the status of the reception, including address match information will be available to the CPU for inspection. The CPU will need to check this field in order to discern what type of MAC frame was just received.

- A second type of MAC traffic could be seen by the repeater bridge. A
- 25           second example of Host traffic is due to the support of the Spanning Tree protocol. Whenever bridges running the Spanning Tree need to communicate with one another, a Bridge Protocol Data Unit (BPDU) is used. Bridges will use BPDUs to communicate changes in the LAN topology, thereby causing a re-evaluation of the best routes a data packet can take. Bridge protocol data units
- 30           are used to help a bridge decide which port offers the shortest path back to the Spanning Tree's root bridge.

The BPDU is addressed to an 802.1d group address. That is, it does not uniquely identify a port, but instead it identifies a unique group; i.e., all 802.1d compliant bridges. Therefore, the frequency of receiving Host traffic could be non-trivial if there are a lot of bridges around the network. As a result, the  
5 repeater may support a third DMA channel which can transfer Host frames to a ring within a different region of memory. A ring structure is provided so that the CPU could be given enough time to process the buffered Host frames.

Regardless of the type of Host traffic, the repeater bridge must be able to source packets from memory. When the software has completed building a  
10 packet in the Host traffic region of memory, it gives a transmit indication to the bus master through a slave write to the control register. This ability to source Host traffic to any port creates an effective fourth DMA channel.

Each companion chip must share its current pointers with the other. Communication between bus masters may be provided using a semaphoring  
15 scheme as disclosed in co-pending patent application, Serial No. 08/528,112, entitled METHOD AND APPARATUS FOR IMPLEMENTING A TWO-PORT ETHERNET BRIDGE USING A SEMAPHORING TECHNIQUE, which is incorporated by reference herein. A receiving media access controller receives a data packet on a first port. The receiving Media Access Controller places the  
20 packet into structure in memory (by performing an access cycle on the bus). The receiving Media Access Controller then indicates the presence of that packet to the transmitting Media Access Controllers by executing a bus cycle to one of the first unique address(es). Subsequently the transmitting Media Access Controller device transmits that packet to users attached to its ports. The transmitting  
25 Media Access Controller then signals back to the receiving Media Access Controller that the packet has been transmitted and that the space in memory which it occupied is now available, by executing a bus cycle to the other (second) unique address. The receiving Media Access Controller is now free to use that space for new packets that it may receive.

30 The two unique addresses that are used by the Media Access Controllers to communicate with each other are programmed into registers on each device. By building the memory structure as a ring, and by defining a register known as

the "Write Boundary" register the two Media Access Controllers are able to pass packets back and forth indefinitely. Each new packet is put into the ring just after the previously received packet. The Write Boundary register indicates the point in the ring past which packets previously received by the first Media

5 Access Controller have not yet been transmitted by the second Media Access Controller. Whenever the transmitting Media Access Controller signals that it has transmitted a packet back to the receiving Media Access Controller, it also updates the Write Boundary register of the receiving Media Access Controller. This is done by transmitting the value of the Write Boundary register during the

10 data portion of the bus cycle.

The foregoing description of the preferred embodiment of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It

15 is intended that the scope of the invention be limited not with this detailed description, but rather by the claims appended hereto.

## WHAT IS CLAIMED IS:

1. A repeater having a plurality of ports for providing connections between media segments coupled thereto, comprising:
  - ten megabit per second repeater logic for connecting ten megabit per
  - 5 second devices to a ten megabit per second backplane;
  - one hundred megabit per second repeater logic for connecting one hundred megabit per second devices to a one hundred megabit per second backplane; and
  - a port switching fabric, coupled to the ten megabit per second repeater
  - 10 logic and to the one hundred megabit per second repeater logic, for determining the transmission rate of a signal at a port and routing the signal to one of the repeater logic according to the transmission rate determination.
2. The repeater of claim 1 further comprising an integrated serial
- 15 controller for accessing internal management counters, control and status configuration information.
3. The repeater of claim 1 further comprising an LED blink scheme for providing status information about the repeater.
- 20
4. The repeater of claim 1 further comprising a media access controller for controlling functions including the transmission, loopback and reception of Ethernet frames.
- 25
5. The repeater of claim 4 further comprising direct memory access controller for servicing the media access controller and controlling memory.
6. The repeater of claim 5 further comprising a transmit and receive FIFO, coupled to the MAC, for providing latency for the DMA controller.
- 30
7. The repeater of claim 4 wherein the media access controller further comprises a receiver for receiving a serial bit stream of data through a

port, stripping off the preamble/SFD header from the data stream, and deserializing the data stream into word quantities.

8. The repeater of claim 5 wherein the FIFO buffer receives the  
5 deserialized data from the media access controller before service by the direct memory access controller.

9. The repeater of claim 5 wherein the direct memory access  
controller transfers data from the receive FIFO to memory during a reception,  
10 and from memory to a transmit FIFO during transmission.

10. The repeater of claim 9 further comprising a dual buffer ring structure for managing the transmit and receive traffic.

11. The repeater of claim 10 wherein the dual buffer ring structure  
15 further comprises two static registers and two dynamic registers.

12. The repeater of claim 11 wherein the registers are programmed so  
20 that the ring boundaries are on 2048 byte pages and the data buffers start on 8 byte boundaries.

13. The repeater of claim 11 further comprising a buffer area disposed  
between the two rings in memory to prevent packet fragmentation so that an  
25 arriving packet that does not fit within the space allotted between the buffer start address and the ring stop boundary is written past the ring stop boundary until the data buffer is complete and the start of the next frame will be wrapped to the ring's start address.

14. The repeater of claim 1 further comprising bridging means having  
30 a receiving and forwarding state, the receiving state for receiving signals from

remote media access control managers and the forwarding state for responding to remote media access control managers.

- 15           15.     A repeater having a plurality of ports for providing connections  
5     between media segments coupled thereto, comprising:  
            a plurality of backplanes operating at different transmission rates;  
            a plurality of repeater logic means, coupled to the plurality of backplanes,  
for connecting a device to at least one of the backplanes;  
            a port switching fabric, coupled to the plurality of repeater logic means,  
10     for determining the transmission rate of the device at a port and routing a signal  
from the device to a repeater logic means according to the transmission rate  
determination; and  
            bridging means, coupled to the repeater logic means, for bridging signals  
between the repeater logic means.

15

            16.     The repeater of claim 15 further comprising a serial controller for  
accessing internal management counters, control and status configuration  
information.

20

            17.     The repeater of claim 18 further comprising LED for providing  
status information about the repeater.

            18.     The repeater of claim 15 further comprising a media access  
controller for controlling functions including the transmission, loopback and  
25     reception of Ethernet frames.

            19.     The repeater of claim 18 further comprising direct memory access  
controller for servicing the media access controller and controlling memory.

            20.     The repeater of claim 19 further comprising a transmit and receive  
30     FIFO, coupled to the MAC, for providing latency for the direct memory access  
controller.



21. The repeater of claim 18 wherein the media access controller further comprises a receiver for receiving a serial bit stream of data through a port, stripping off the preamble/SFD header from the data stream, and deserializing the data stream into word quantities.

5

22. The repeater of claim 19 wherein the FIFO buffer receives the deserialized data from the media access controller before service by the direct memory access controller.

10

23. The repeater of claim 19 wherein the direct memory access controller transfers data from the receive FIFO to memory during a reception, and from memory to a transmit FIFO during transmission.

15

24. The repeater of claim 23 further comprising a dual buffer ring structure for managing the transmit and receive traffic.


25. The repeater of claim 24 wherein the dual buffer ring structure further comprises two static registers and two dynamic registers.


20

26. The repeater of claim 25 wherein the registers are programmed so that the ring boundaries are on 2048 byte pages and the data buffers start on 8 byte boundaries.

25

27. The repeater of claim 25 further comprising a buffer area disposed between the two rings in memory to prevent packet fragmentation so that an arriving packet that does not fit within the space allotted between the buffer start address and the ring stop boundary is written past the ring stop boundary until the data buffer is complete and the start of the next frame will be wrapped to the ring's start address.





**FIG. 2**




FIG. 3




FIG. 4




FIG. 5




FIG. 6

7/9




FIG. 7






FIG. 8

9/9




FIG. 9

# INTERNATIONAL SEARCH REPORT

International Application No  
PC1/US 96/14197

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 6 H04L12/44

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO,A,91 16679 (CHIPCOM, CORP.) 31 October 1991  see page 2, line 26 - page 3, line 15 see page 5, line 4 - line 13 see page 15, line 2 - line 17 see page 19, line 23 - page 20, line 24 see page 48, line 1 - line 18 see page 49, line 8 - line 19 see figures 2,3	1,2,4-9, 14-16, 18-23
A	WO,A,95 34976 (NATIONAL SEMICONDUCTOR CORPORATION) 21 December 1995  see page 1, line 26 - line 35 see page 2, line 22 - line 27 see page 3, line 2 - page 4, line 24 see page 5, line 29 - page 6, line 25  -/-	1,2,4,7, 15,16, 18,21

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

\* Special categories of cited documents :

- \*A\* document defining the general state of the art which is not considered to be of particular relevance
- \*E\* earlier document but published on or after the international filing date
- \*L\* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- \*O\* document referring to an oral disclosure, use, exhibition or other means
- \*P\* document published prior to the international filing date but later than the priority date claimed

\*T\* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

\*X\* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

\*Y\* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

\*&\* document member of the same patent family

Date of the actual completion of the international search

18 December 1996

Date of mailing of the international search report

02.01.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax (+31-70) 340-3016

Authorized officer

Vaskimo, K

# INTERNATIONAL SEARCH REPORT

International Application No  
PCI/US 96/14197

## C(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>ELECTRONIC DESIGN, vol. 43, no. 6, 20 March 1995, CLEVELAND, OH, US, pages 155-156, 158, 160, XP000509380 LEE GOLDBERG: "100Base-T4 Transceiver Simplifies Adapter, Repeater, And Switch Designs" see page 155, line 1 - line 34 see page 158, left-hand column, line 5 - line 16 see page 160, left-hand column, line 50 - right-hand column, line 30 see figures 1,2</p> <p style="text-align: center;">---</p>	<p>1,2,4-9, 15,16, 18-23</p>
A	<p>DATA COMMUNICATIONS, vol. 21, no. 17, 21 November 1992, NEW YORK, US, pages 56-58, 60-62, 64, 66, 68, XP000325098 KEVIN TOLLY: "Grading Smart Hubs for Corporate Networking" see page 58, left-hand column, line 22 - right-hand column, line 43 see page 60, middle column, line 2 - line 10 see page 61, right-hand column, line 33 - line 49</p> <p style="text-align: center;">---</p>	<p>1-4,7, 15-18,21</p>
A	<p>HEWLETT-PACKARD JOURNAL, vol. 46, no. 4, August 1985, PALO ALTO, CA, US, pages 6-12, XP000525580 ALAN R. ALBRECHT, PATRICIA A. THALER: "Introduction to 100VG-AnyLAN and the IEEE 802.12 Local Area Network Standard" see page 10, left-hand column, line 17 - page 11, left-hand column, line 7 see figures 4,5</p> <p style="text-align: center;">-----</p>	<p>1,15</p>

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 96/14197

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO-A-9116679	31-10-91	US-A- 5301303	05-04-94
		CA-A- 2079671	24-10-91
		EP-A- 0526580	10-02-93
		US-A- 5530842	25-06-96
		IL-A- 98771	27-02-94
-----			
WO-A-9534976	21-12-95	US-A- 5541957	30-07-96
		EP-A- 0713623	29-05-96
-----			

**THIS PAGE BLANK (USPTO)**